

Next-Generation Application Delivery

yankee
group

www.yankeegroup.com

by Zeus Kerravala | April 2007

Executive Summary

The network plays a critical role in the availability, performance and security of enterprise applications. If the network is not properly deployed, the impact can be severe. According to respondents of the Yankee Group 2005 Application Management Survey, application performance issues result in an average productivity decrease of 14% (see Exhibit 1). This is evidence that application performance—as well as availability—directly affects employee productivity and the bottom line of a company. As more users work more hours leveraging key business applications, it becomes even more important to address application availability and performance issues to ensure achievement of business processes and objectives.

Application availability and acceleration technologies have evolved during the past several years. Technologies exist to enhance application availability, security and performance within the data center—at the branch and within the WAN itself. Network managers need to determine the right mix of data center and branch application delivery technologies, as well as the most effective architecture for today and the future.

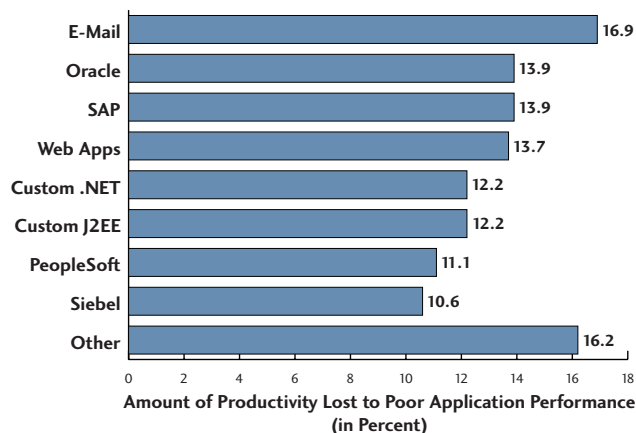
Most organizations have already deployed some application availability or acceleration technologies as overlays to their existing network. Most products were deployed ad hoc to address one or two specific application delivery issues: server load and availability, or mitigation of WAN latency and poor application response times. With nearly every application now being made available over the WAN, it's time for enterprises to integrate true application fluency—maximizing availability, security and performance—into the network.

An integrated, application-fluent approach is the best way to leverage the network to deliver all applications, optimize network and infrastructure investments and visibly reduce operational and capital expenses.

Exhibit 1

The Impact of Poor Application Performance

Source: Yankee Group 2005 Application Management Survey



Note: The average for all applications is 14%.
n = 252

This custom publication has been funded by Cisco Systems.

© Copyright 1997-2007. Yankee Group Research, Inc. All rights reserved.

This Yankee Group Consulting Report is published for the sole use of Yankee Group clients. It may not be duplicated, reproduced or transmitted in whole or in part without the express permission of Yankee Group, 31 St. James Avenue, Boston, MA 02116. For more information, contact Yankee Group: info@yankeegroup.com; Phone: 617-956-5005. All rights reserved. All opinions and estimates herein constitute our judgment as of this date and are subject to change without notice.

Table of Contents

| | | |
|------|--|----|
| I. | Enterprise Application Delivery Challenges | 3 |
| II. | Application Infrastructure Requirements and Evolution | 4 |
| | Unique Application Requirements | 4 |
| | Application Networking Equipment Evolution | 5 |
| III. | The Vision for Application Networking Equipment | 6 |
| | Application Infrastructure Overlay | 6 |
| | Integrated, Application-Fluent Network Infrastructure | 7 |
| IV. | Cisco's Application Networking Offering | 8 |
| | Cisco ANS Delivers an End-to-End, Application-Fluent Network | 9 |
| | Cisco ANS: The Value of an Integrated vs. Overlay Architecture | 10 |
| V. | Choosing the Best Solution for an Organization | 10 |
| VI. | Conclusions | 11 |

I. Enterprise Application Delivery Challenges

More than ever, organizations rely on network-delivered client/server applications. According to Yankee Group research and multiple industry studies, approximately 80% of the enterprise workforce works outside of company headquarters. For IT departments, the challenge is to deliver consistent application availability and performance—coupled with the required levels of security—whether these users are at corporate headquarters or are accessing the application over a wide-area network (WAN) from a branch office or a remote location.

The difficulty is that most enterprise applications were not written or designed for deployment outside of corporate headquarters. They were written years ago for deployment on a local-area network (LAN), where there is plenty of bandwidth, minimal latency and a limited number of users. As a result, users accessing the applications from branch or traveling locations often experience poor performance, which can vary widely by location or time of day.

An additional challenge for IT departments is managing the ever-growing quantity of applications to deliver, while maintaining performance and limiting the quantity of related infrastructure within the data center.

In addition, application and data center security risks increase significantly as more applications are deployed and accessed over the WAN by an expanding number of end users.

There are many factors that make applications difficult to deploy and deliver effectively over the network. The most common issues include:

- **Inflexible application infrastructure:** Application design has historically been done on an application-by-application basis. This means that the infrastructure used for a particular application is often unique to that application. This type of design tightly couples the application to the infrastructure and offers little flexibility. Consequently, it takes longer to deploy new applications, and the cost is much higher. Because the

application and infrastructure are tightly coupled, it is difficult to partition resources and levels of control to match business requirements. IT is forced to deploy dedicated infrastructure that's underutilized and difficult to consolidate.

- **Server availability and load:** As applications are deployed enterprise-wide and access to them increases to partners and customers, the server-processing load related to networking tasks continues to increase year over year. Network- and inter-application-related processes such as SSL and XML/SOAP traffic usage are increasing, and the servers are expected to process a higher load. The mission-critical nature of applications puts a premium on server availability. Despite the benefits of server virtualization technology, the number of physical servers continues to grow based on new application deployments, raising power and cooling requirements.
- **Application infrastructure capacity:** Enterprise data centers and WAN infrastructure are designed to deliver critical applications for an increasing number of end users, whether these are local, mobile or branch office workers. IT organizations are seeing exponential growth in transaction volume, while simultaneously introducing converged services such as voice and video over IP. New software capabilities and wider, more distributed deployment often exceed the capacity of existing infrastructure.
- **Adoption of next-generation WANs:** Enterprises have migrated from dedicated point-to-point and frame relay networks and widely adopted internet and VPN technologies. Next-generation WANs are much more cost-effective, but they may also introduce latency, jitter and packet loss, which may have a significant effect on certain applications that are delay-sensitive. In addition, many enterprises deploy multiple point appliances to address application performance and security issues. The result is a more complex and costly application infrastructure as well as increased latency associated with traversing multiple appliances.

- **Application security and compliance requirements:** Security is not just about preventing unauthorized access, providing secure transport of application data or ensuring server and network availability. Many of the new security events are the result of application- and document-embedded attacks that compromise application performance and availability. Such attacks also potentially cause loss of vital application data—even while leaving networks and servers unaffected. For example, deliberately malformed application requests can reach server banks and significantly downgrade performance or make applications unavailable. In addition, attacks can spread between application tiers and cause widespread problems.

One possible solution to improve application performance and availability is to rewrite the application completely to make it network-optimized. However, this requires application developers to have a much deeper understanding of how different applications respond to things such as bandwidth constraints, delay, jitter and other network variances. In addition, developers need a clearly predictable view of an end user's foreseeable access method. This is simply not feasible for every enterprise application—particularly legacy applications that took years to write and customize and that can contain millions of lines of code.

In today's complex enterprise environment with a large number of applications being delivered to a large number of users at multiple locations connected across a shared network, the more effective solution is to embed application fluency and enablement within as much of the network infrastructure as possible.

Application fluency gives the network visibility into each application flow, its unique performance and security requirements and real-time application traffic interdependencies. The ability of the network to understand applications enables the network to “tune” itself automatically to optimize application performance, availability and security. In short, application fluency makes the application successful in achieving its IT and business objectives.

To understand how the network can automatically tune itself based on the applications it supports, it's necessary to understand how application network infrastructure has evolved. Many point solutions for particular application delivery problems have already become features or techniques within the network infrastructure. The best strategy for improving application availability and performance is a comprehensive solution that has true application fluency and works to optimize applications from the data center to the branch.

II. Application Infrastructure Requirements and Evolution

Unique Application Requirements

Different types of applications respond differently to the many application delivery and acceleration methods available today. One method can help one application while degrading the performance of another. The primary types of applications and their most effective delivery techniques are:

- **General TCP applications:** TCP-based client/server applications are typically not optimized for the latency, bandwidth and congestion characteristics of WANs. Latency-associated application delays are typically visible because networking protocols (e.g., TCP/IP) and application protocols (e.g., file access protocols such as CIFS/NFS) were not designed for the WAN. For example, when a file is opened, there are numerous “handshakes” and acknowledgements sent between the client and server. This is why TCP/IP is often referred to as a chatty protocol. In addition, many client/server applications share data across users and user sessions. TCP-based client/server applications typically respond well to WAN optimization techniques such as data suppression and compression, application protocol and TCP latency optimization and wide-area file services (WAFS).

- **Web-based applications:** Web-based applications are best optimized using a combination of techniques, including request offloading and connection multiplexing in the data center and transport optimization across the WAN. Together, these techniques can mitigate the server load and the chattiness often experienced with web applications.
- **XML and service-oriented applications:** Many newer applications include next-generation protocols such as XML to enable communication between different application components or service modules. The result is increased server load and security requirements, both of which network-based XML offload and security technologies can address.
- **UDP applications:** Real-time UDP applications cannot be accelerated, although they can be optimized with QoS. The challenge is optimizing QoS policies and UDP applications without imposing performance degradation on other applications.
- **Video and digital media applications:** The use of video and other digital media for applications such as e-learning, corporate communications and digital signage is becoming increasingly more popular in organizations. The network is expected to optimize the various types of media traffic to ensure that high-quality, long-playing videos can be streamed live, delivered on demand or multicast across the organization. This will result in improved user experiences at a lower cost.

The challenge for network and application teams is to balance the different application delivery and acceleration techniques for their unique environment. For many applications, front-end application acceleration can greatly improve response time. In other cases, off-loading CPU-intensive processes, such as SSL, from the servers in the data center can best improve performance. Additionally, as applications become more modular and service-oriented architecture (SOA) adoption increases, the amount of XML traffic will definitely rise.

Therefore, it is critical to select vendors that can provide a broad range of application-fluent networking technologies to ensure investment protection as well as ongoing interoperability with existing and future applications.

Application Networking Equipment Evolution

Application networking equipment integrates the network infrastructure with applications and application infrastructure. The more integrated the network and application infrastructures, the more effectively the network can deliver applications to the growing end-user base.

Application networking equipment has evolved in two primary areas of deployment: the data center and the branch.

Data Center Application Networking

The idea of improving application performance began in the data center. The internet boom ushered in the era of the server load balancers (SLB). SLBs balance the load on server banks to improve their response to client requests, although they have evolved and taken on additional responsibilities such as application proxies and complete Layer 4 through 7 application switching.

From their mainstream roles in providing Layer 4 through 7 switching, today's application switches now provide an array of acceleration and server offload benefits, including TCP processing offload, SSL offload, compression and various other acceleration technologies. An emerging category of devices provides additional availability and security benefits for XML-based applications: offloading the processing of XML transactions and inspecting traffic for incremental security.

These new devices sit within the enterprise data center in front of the web and application servers and provide a range of services to maximize server availability, security and asymmetric (from server to client browser) application acceleration.

Branch Application Networking (WAN Optimization)

It is not uncommon to see branch offices maintaining their own mini IT department—a dedicated IT infrastructure of servers, storage and applications that often requires local staff to manage. But this practice can become very expensive and has led many organizations to consolidate branch office infrastructure and virtually deliver applications and services over the WAN from central data centers. To handle the corresponding increase in WAN traffic as well as the impact on application and file performance, network and application teams have been incorporating technologies such as QoS, caching, latency mitigation, compression and WAFS—collectively called WAN optimization:

- **QoS:** QoS prioritizes certain classes of traffic over others. Using QoS, network managers can prioritize voice and video traffic over non-real-time applications traffic.
- **Caching:** Caching is simply a way to separate static content from applications and cache it locally so the same content is not being continuously retransmitted.
- **Latency mitigation:** Latency mitigation improves application performance by virtually moving the application closer to the source of the data. The application converses with the WAN optimizer instead of the server across the WAN, shortening the overall round-trip time of the packets.
- **Compression:** Compression involves algorithms that can reduce or eliminate unnecessary packet content transmission. Different compression algorithms will result in different loss levels—more loss usually means faster transmissions. Network managers need to evaluate the loss and determine the best loss/performance ratio for the application.
- **Wide-area file services:** WAFS improve the performance of file access and sharing over a WAN.

Additionally, many TCP applications perform poorly from a branch office because of latency, chattiness or other issues related to the distance between the branch and the headquarters. TCP optimization technology can reduce the effective distance between the branch office and the data center, enabling the application in many cases to perform as if it were running over a LAN rather than over a WAN.

III. The Vision for Application Networking Equipment

Application networking equipment has evolved from a series of point technologies and products, forcing enterprises to now take a more strategic and holistic view. The bottom line is that the network needs application fluency to maximize application availability and performance without increasing complexity and risk.

To accomplish this network need, the choices are to deploy infrastructure as an overlay to the network or to embed application fluency within the network infrastructure.

Application Infrastructure Overlay

Many enterprises, by virtue of having deployed application networking in an ad-hoc way, already have an overlay of application networking equipment. To continue this approach, enterprises deploy dedicated application networking that is connected to, but separate from, the network infrastructure. Like any approach, application infrastructure overlay has some pros and cons.

Pros

The initial reason for deploying dedicated application networking equipment is so network and application teams can quickly deploy a point solution to a specific problem. For example, if a key centralized enterprise application performs poorly over the WAN, then a point WAN optimization solution can quickly address the issue.

In addition, that point product includes its own management tools that provide some quick visibility into what's on the network. That visibility alone provides some value and enables network managers to resolve a specific application performance issue quickly. Dedicated application networking also enables network managers to separate acceleration and optimization technologies from basic network tasks such as routing.

Cons

One of the most significant cons to dedicated application networking is the lack of integration between the application networking equipment and the broader network infrastructure. A standalone application networking product may require changes to router or security configurations; it can affect the optimal settings of the network infrastructure. Network managers then need to adjust the operating system and configurations of the network devices manually—a tall order in a large, distributed network.

In addition, different optimization technologies can affect the performance of other applications. For example, accelerating Microsoft Exchange with an overlay technology could break the QoS running on a branch router, which could affect the performance of VoIP. Balancing the different methods can be a challenge, and not doing so effectively can actually hurt application performance and availability.

Network managers must contend with separate management domains—one for the application networking equipment and another for the network infrastructure. This is on top of an already complex IT operations and management infrastructure. It is difficult to get to the root cause of an application performance issue when troubleshooting requires views into multiple management systems that are completely separate.

Finally, the overlay approach mandates the management and maintenance of two infrastructures. This requires training staff on multiple technologies and increases maintenance costs.

Integrated, Application-Fluent Network Infrastructure

In contrast to the overlay approach, integrating application networking and overall network infrastructure puts application intelligence and fluency within the network gear. There is no separate application network infrastructure; the application fluency is a function of the network. This empowers the network itself to make application-smart decisions. As with application infrastructure overlay, integrated, application-fluent network infrastructure has pros and cons.

Pros

The most important benefit of an integrated, application-fluent network infrastructure is having a single infrastructure for application and network optimization. Integration of the two minimizes or eliminates contention between application availability/acceleration techniques and Layer 2 and 3 routing and switching.

The integrated approach also creates a single infrastructure for enhancing data center and branch optimization technologies. Enterprises can more effectively target consolidation efforts and balance their data center and branch technology. It can also mean more efficient optimization and accelerated performance because there are fewer bumps in the wire.

Having one application and network infrastructure also reduces operational expense (opex) considerably (e.g., sharing advanced application switching, security and basic Layer 3 switching in a single chassis, or having routing, QoS and WAN optimization in one branch device). Deployment, provisioning and management are simpler because it is one infrastructure. The integrated approach also reduces capital expense (capex) because features do not require new hardware devices; network managers can instead add and migrate more easily to new features.

Cons

For some organizations, an integrated strategy can mean removing pieces of an overlay network before receiving a full return on the investment. It may also require some upgrades to existing network equipment.

Network and application teams also need to expand their skills to include the new application-fluent networking and security technologies. However, much of this is offset by the significant reductions that an enterprise can experience in capex and opex.

After measuring the pros and cons of both approaches, Yankee Group believes that integrating application networking within network infrastructure—integrated application fluency—is the best approach. It's inevitable that most application availability, security, acceleration and WAN optimization technologies will evolve to become features of the network. Enterprises need an integrated approach to optimize data center and branch resources when networked applications are the drivers for the business.

IV. Cisco's Application Networking Offering

The Yankee Group 2005 Application Management Survey revealed that 70% of network managers are responsible for actively troubleshooting application performance issues. They can no longer get away with saying, "It's not the network." As the networking market leader, Cisco understands the network's increasing role in application delivery and has introduced during the last year different components of its data center-to-branch application networking solution, Cisco Application Networking Services (ANS).

Cisco ANS is a set of solutions that addresses performance, availability and security issues in the data center, at the branch and within the WAN. Cisco states that the branch and WAN products enable branch office infrastructure consolidation by delivering LAN-like application performance. The data center products optimize and secure both applications and related servers.

Key elements of Cisco's ANS strategy include Cisco Wide Area Application Services (WAAS) for branch and WAN optimization and Cisco Application Control Engine (ACE) for data center optimization and security. WAAS and ACE integrate with Cisco IOS Software to provide data center and application optimization within the network infrastructure and to realize the vision of the application-fluent network—a single network with the intelligence to view deeply into application traffic, make decisions and implement services to optimize application results.

Cisco ANS Delivers an End-to-End, Application-Fluent Network

Cisco is unique among vendors in the application networking equipment market because of the breadth of its solution portfolio—from appliance- and router-integrated WAN optimization, video delivery systems, legacy and next-generation application switches and application firewall functionality, along with application fluency embedded in its range of router and switch products.

Cisco's ACE products for data center application networking give IT departments more control over application and server infrastructure, which enables them to manage and secure application services more easily and improves performance. The ACE product family provides several advantages for network and application teams, including high-performance Layer 4 and 7 switching, high-capacity application firewall capabilities through detailed packet inspection and server offloading of various processing tasks such as SSL and TCP.

A unique capability of the ACE product family is its ability to provide up to 250 virtual devices from a single platform, thereby reducing capex and power and cooling costs (target ROI can be less than 12 months), which results in less complexity and lower overall TCO. Additionally, the role-based administration capability enables multiple network and application teams to configure their respective portions of the device in parallel, enabling faster application deployments.

WAAS is the key enabler of Cisco's vision for application performance and fluency within a consolidated branch infrastructure. It includes a wide range of application acceleration, WAN optimization and WAFS technologies. WAAS seamlessly integrates with the network infrastructure and can be deployed as an appliance (wide-area application engine or WAE) or as a module for the integrated services router (ISR).

The combination of Cisco IOS/WAAS technologies addresses latency issues, bandwidth management, link throughput and traffic prioritization; it can replace local branch services with services delivered from a central data center. Multiple point solutions from various vendors can address many of these application performance issues. However, what distinguishes Cisco WAAS is its tight integration with IOS and security environments, transparently leveraging existing services such as NetFlow, quality of service and firewall access control lists (ACLs). As a result, Cisco WAAS preserves operating network services and reduces ongoing opex.

Cisco WAAS understands the application traffic and can apply the technology that best suits that traffic. Specifically, Cisco application adapters are proxies that intercept application and protocol-level messages and handle them locally. This eliminates a lot of network chatter. It also reduces transmissions over the WAN while maintaining normal application behavior and data integrity. In addition, WAAS uses a combination of data suppression and compression techniques to significantly reduce the amount of redundant data that is repeatedly transmitted across the WAN as well as TCP optimization techniques to reduce the delays associated with TCP transmission across the wire. In short, WAAS understands the objects that applications use and handles them more efficiently.

Cisco ANS: The Value of an Integrated vs. Overlay Architecture

Most distributed enterprises have elements of an overlay application infrastructure in place. However, most organizations addressed data center, WAN and application optimization separately. The result is a separately managed infrastructure that is difficult to optimize for changing application requirements. Most importantly, an overlay application infrastructure is not integrated with the network infrastructure. As a result, it can be intrusive and can require additional management and integration tasks to avoid service disruption and performance issues.

In an overlay architecture, traditional WAN optimization changes TCP/IP header information. It creates an optimization tunnel between a core and an edge device and applies acceleration techniques. But by changing the TCP/IP header information, some services may not work and network staff will need to spend more time integrating the WAN optimization with the network infrastructure. The dedicated link can also create a single point of failure. In addition, when network managers make changes to the WAN optimization devices, it often requires that they make configuration changes to the network infrastructure as well. The result is an application infrastructure that is less efficient that can be potentially harmful as well as increasing equipment and operational costs.

In contrast, an integrated approach preserves TCP/IP header information, so services are not affected. WAN optimization and application acceleration techniques are integrated with network features such as Cisco IOS and AutoQoS. Network managers manage one platform for application and network infrastructure, which reduces operational costs. Management is more seamless. Cisco IP service-level agreements (SLAs) provide instrumentation and performance reporting for both the network and application infrastructures, so network managers have an integrated view of the network and application infrastructure. Changes and the impact of changes are easier to see and manage.

An integrated approach also provides benefits within the data center, where Cisco's ACE application switch integrates with the Catalyst switch's network and security services, such as campus-wide QoS via Network Based Application Recognition (NBAR). An example of full network integration, ACE can also load balance WAAS devices at the headend to increase WAN optimization capacity, rather than requiring extra point products from a WAN optimization vendor.

By integrating everything from the data center to the branch, enterprises can deploy and manage one set of platforms and empower the network to automate the application of performance optimization techniques for every kind of enterprise application. Operating one integrated application and network infrastructure means enterprises can deploy new applications over the WAN more quickly, better understand the performance and have much lower deployment and operational costs.

V. Choosing the Best Solution for an Organization

Because many organizations have some dedicated load balancing, application acceleration or WAN optimization products in place, it is important to evaluate their current environment, staff skills and operating costs. When evaluating a solution, key questions to ask or include in a request for proposal (RFP) are:

- **What impact (positive and negative) will the solution have on existing network services?** This question will help determine whether an integrated or overlay architecture is best. A solution's value diminishes if it requires additional integration and management.
- **Which application types does the solution optimize best?** This question helps organizations determine if there is a match between the types of applications they need to optimize and what the solution does best. It's also important to know if the solution may negatively affect certain applications (e.g., VoIP).

- **How will the solution support future application deployment?** It's important to understand the application deployment and performance issues that require immediate attention, but it's also crucial to understand how application needs will change and whether the solution can adapt and scale accordingly. XML traffic is on the rise, so organizations should ensure that the vendor of their choice has included it in its product.
- **How much of a performance difference affects the purchase decision?** This question is straight cost/benefit analysis. What is the performance difference among the various solutions, and how much should it be to affect an organization's purchase decision? If a point product delivers better performance, is that offset by the capex and opex savings of an integrated application networking solution?
- **Does the branch and remote user solution integrate with the data center products?** Integration with the data center and WAN and branch solutions is important and will only increase in importance as enterprise seek to deploy more applications over the WAN. Integration makes optimization more efficient, enables more automation and reduces management complexity.
- **How difficult/simple will network integration and ongoing management be?** Organizations should seek solutions that offer an integrated view of network, data center and branch optimization technologies, as well as visibility throughout the network and into application traffic streams.
- **How should organizations leverage advance/professional services to ensure a smooth deployment?** Organizations should evaluate their staff's skills and the solution provider's ability to fill in gaps, transfer knowledge and be available post-sale to ensure successful deployment.

These questions offer a good starting point for assessing an organization's needs. Point products can address immediate problems and answer specific questions. A proper evaluation will equally weigh all the questions and vendors' responses. As managers look to the future of application delivery for their distributed organizations, a solution—not a product—will deliver more value.

VI. Conclusions

Application delivery in distributed enterprises is critical to maintain a competitive advantage. CIOs are under constant pressure to improve the productivity and value of core business applications and simultaneously control costs. To deliver more value, enterprises continue to deploy more mission-critical applications to employees over their IP infrastructure. Point products can help with specific application delivery issues, but companies should increasingly turn to complete application networking solutions that leverage their data center and WAN infrastructure to extend their market reach from their central infrastructure.

Integrating application networking and network infrastructure helps enterprises be certain that end users receive the performance and availability required to ensure application usage and ROI. The future of dynamic application delivery requires that enterprise networks have application fluency and the ability to apply the most effective performance enhancements automatically across the network.

Enterprises looking to enhance the performance and availability of their applications and to reduce operational costs and capital expenditures should look to an integrated application and network infrastructure.

It is our view that Cisco provides a complete and highly effective solution for end-to-end application networking. Its unique advantage is its high level of integration with existing network and application infrastructure and services, coupled with a highly efficient operational model, application fluency across multiple layers and a high level of extensibility.

Yankee Group

Yankee Group has research and sales staff located in North America, Europe, the Middle East, Africa, Latin America and Asia-Pacific. For more information, please contact one of the sales offices listed below.

Corporate Headquarters

31 St. James Avenue
BOSTON, MASSACHUSETTS 02116-4114
617-956-5000 phone
617-956-5005 fax
info@yankeegroup.com

Europe

55 Russell Square
LONDON WC1B 4HP
UNITED KINGDOM
44-20-7307-1050 phone
44-20-7323-3747 fax
euroinfo@yankeegroup.com

Yankee Group | the global connectivity experts™

A global connectivity revolution is under way, transforming the way that businesses and consumers interact beyond anything we have experienced to date. The stakes are high, and there are new needs to be met while power shifts among traditional and new market entrants. Advice about technology change is everywhere—in the clamor of the media, the boardroom approaches of management consultants and the technology research community. Among these sources, Yankee Group stands out as the original and most respected source of deep insight and counsel for the builders, operators and users of connectivity solutions.

For 37 years, we have conducted primary research on the fundamental questions that chart the pace and nature of technology changes on networks, consumers and enterprises. Coupling professional expertise in communications development and deployment with hundreds of interviews and tens of thousands of data points each year, we provide qualitative and quantitative information to our clients in an insightful, timely, flexible and economic offering.

Yankee Group Link

As technology connects more people, places and things, players must confront challenging questions to benefit from the changes: which technologies, what economic models, which partners and what offerings? Yankee Group Link™ is the research membership uniquely positioned to bring you the focus, the depth, the history and the flexibility you need to answer these questions.

Yankee Group Link membership connects you to our qualitative analysis of the technologies, services and industries we assess in our research agenda charting global connectivity change. It also connects you to unique quantitative data from the dozens of annual surveys we conduct with thousands of enterprises and consumers, along with market adoption data, comprehensive forecasts and global regulatory dashboards.

Yankee Group Link Research

As a Link member, you have access to more than 500 research reports and notes that Yankee Group publishes each year. Link Research examines current business issues with a unique combination of knowledge and services. We explore topics in an easy-to-read, solutions-oriented format. With the combination of market-driven research and built-in direct access to Yankee Group analysts, you benefit from the interpretation and application of our research to your individual business requirements.

Yankee Group Link Interaction

Our analysts are at your further disposal with data, information or advice on a particular topic at the core of a Link membership. We encourage you to have direct interaction with analysts through ongoing conversations, conference calls and briefings.

Yankee Group Link Data

Yankee Group Link Data modules provide a comprehensive, quantitative perspective of global connectivity markets, technologies and the competitive landscape. Together with Link Research, data modules connect you to the information you need to make the most informed strategic and tactical business decisions.

Yankee Group Consulting

Who better than Yankee Group to help you define key global connectivity strategies, scope major technology initiatives and determine your organization's readiness to undertake them, differentiate yourself competitively or guide initiatives around connectivity change? Our analysts apply Yankee Group research, methodologies, critical thinking and survey results to your specific needs to produce expert, timely, custom results.

Yankee Group Live!

The global connectivity revolution won't wait. Join our live debates to discuss the impact ubiquitous connectivity will have on your future. Yankee Group's signature events—conferences, webinars and speaking engagements—offer our clients new insight, knowledge and expertise to better understand and overcome the obstacles to succeed in this connectivity revolution.

www.yankeegroup.com

The people of Yankee Group are the global connectivity experts™—the leading source of insight and counsel for builders, operators and users of connectivity solutions. For more than 35 years, Yankee Group has conducted primary research that charts the pace of technology change and its effect on networks, consumers and enterprises. Headquartered in Boston, Yankee Group has a global presence including operations in North America, Europe, the Middle East, Africa, Latin America and Asia-Pacific.